arXiv:1611.00347v1 [math.OC] 1 Nov 2016

SURPASSING GRADIENT DESCENT PROVABLY: A CYCLIC INCREMENTAL METHOD WITH LINEAR CONVERGENCE RATE*

ARYAN MOKHTARI†, MERT GÜRBÜZBALABAN‡, AND ALEJANDRO RIBEIRO†

Abstract. Recently, there has been growing interest in developing optimization methods for solving large-scale machine learning problems. Most of these problems boil down to the problem of minimizing an average of a finite set of smooth and strongly convex functions where the number of functions n is large. Gradient descent method (GD) is successful in minimizing convex problems at a fast linear rate; however, it is not applicable to the considered large-scale optimization setting because of the high computational complexity. Incremental methods resolve this drawback of gradient methods by replacing the required gradient for the descent direction with an incremental gradient approximation. They operate by evaluating one gradient per iteration and executing the average of the n available gradients as a gradient approximate. Although, incremental methods reduce the computational cost of GD, their convergence rates do not justify their advantage relative to GD in terms of the total number of gradient evaluations until convergence. In this paper, we introduce a Double Incremental Aggregated Gradient method (DIAG) that computes the gradient of only one function at each iteration, which is chosen based on a cyclic scheme, and uses the aggregated average gradient of all the functions to approximate the full gradient. The iterates of the proposed DIAG method uses averages of both iterates and gradients in oppose to classic incremental methods that utilize gradient averages but do *not* utilize iterate averages. We prove that not only the proposed DIAG method converges linearly to the optimal solution, but also its linear convergence factor justifies the advantage of incremental methods on GD. In particular, we prove that the worst case performance of DIAG is better than the worst case performance of GD. Numerical experiments on quadratic programming and logistic regression problems showcase the advantage of DIAG relative to GD and other incremental methods.

Key words. Incremental methods, finite sum minimization, large-scale optimization, linear convergence rate, worst case analysis

AMS subject classifications. 90C06, 90C25, 90C30, 90C52

1. Introduction.

This paper focuses on finite sum optimization where the objective function can be written as the sum of a set of strongly convex functions. In particular, consider $\mathbf{x} \in \mathbb{R}^p$ as the optimization variable and $f_i : \mathbb{R}^p \to \mathbb{R}$ as the i-th available function. We aim to find the minimizer of the average function $f(\mathbf{x}) = (1/n)\sum_{i=1}^n f_i(\mathbf{x})$, i.e., we intend to solve the optimization problem

$$\mathbf{x}^* = \operatorname*{argmin}_{\mathbf{x}\in\mathbb{R}^p} f(\mathbf{x}) := \operatorname*{argmin}_{\mathbf{x}\in\mathbb{R}^p} \frac{1}{n}\sum_{i=1}^n f_i(\mathbf{x}). \tag{1}$$

In this paper, we refer to f_i as the instantaneous functions and the average function f as the global objective function. This class of optimization problems arises in many fields such as machine learning [5, 4, 25, 8], optimal control, [6, 7, 16], and wireless communications [20, 21]. Our focus is on problems where the instantaneous functions f_i are smooth and strongly convex.

To explain the contribution of this paper we have to discuss the rate and constants that characterize convergence of the different first order methods that can be used to solve the problem in (1). To begin with we can neglect the specific form of f and use the conventional gradient descent (GD) method which is known to converge linearly to

*This work was supported by NSF CAREER CCF-0952867 and ONR N00014-12-1-0997.

†Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA (aryanm@seas.upenn.edu,aribeiro@seas.upenn.edu)

‡Department of Management Science and Information Systems, Rutgers University, Piscataway, NJ (mgurbuzbalaban@business.rutgers.edu).

the optimal argument [19]. This linear convergence rate comes from using individual iterations that are very costly when the number of functions n is large and motivates the use of stochastic and incremental methods in which only one of the instantaneous gradients ∇f_i is evaluated at each iteration. The selection is random in stochastic methods and cyclic in incremental methods. In either case the idea is that individual iterations are less efficient but since n stochastic or incremental operations have the same cost as one GD iteration, overall convergence is faster.

Although faster convergence is observed in many practical situations, it is not known if it is possible to design a stochastic or incremental method with convergence guarantees that are better than the convergence guarantees of GD. The stochastic gradient descent (SGD) method, for instance, is known to have a sublinear convergence rate and is therefore surpassed by regular GD as the number of iterations grows [22]. This limitation is in fact the motivation for alternative stochastic descent methods that achieve linear convergence rates by reducing the variance of stochastic descent directions. Examples of this growing and consequential literature includes stochastic averaging gradient algorithms [23, 9, 10, 18], variance reduction methods [13, 31], dual coordinate methods [26, 27], hybrid algorithms [32, 14], and majorization-minimization algorithms [17]. All of these stochastic methods are successful in achieving a linear convergence rate in expectation with individual iterations that have cost comparable to the cost of SGD iterations. However, the linear convergence constants of these methods are not necessarily better than the linear convergence constant of GD for a problem with comparable condition number. This leaves open the possibility that the worst case performance of these methods is worse than the worst case performance of GD – see Section 2.

Given that the only difference between stochastic and incremental methods is that in the latter functions are chosen in a *cyclic* order – as opposed from the selection in stochastic methods which is uniformly at random – it is not surprising that analogous statements can be made for incremental gradient descent methods (IGD) [1, 28, 3, 2, 29, 11, 30]. Standard IGD has a slow sublinear convergence rate, which motivates the introduction of memory. This is done in the definition of the incremental aggregated gradient (IAG) method that is shown to achieve linear convergence [11] but with a constant that is not necessarily better that the GD constant. Thus, and as in the case of stochastic methods, it is possible that the worst case performance of IAG is worse than the worst case performance of GD – see Section 2.

The contribution of this paper is to introduce a first order incremental method that has a linear convergence rate with a constant that is better than the GD constant of a problem with comparable condition number. This means that the worst case performance of the proposed algorithm is guaranteed to be no worse than the worst case performance of GD. The algorithm relies on keeping memory of past variable *and* gradient evaluations and is therefore termed the Double Incremental Aggregated Gradient (DIAG) method to emphasize the difference with regular IAG methods in which only gradient histories are maintained. This major difference comes from the fact that DIAG uses a different approximation of the global function f at each iteration from the one used in IAG. In particular, DIAG approximates each instantaneous function f_i by the sum of its first order approximation and a proximity term, both evaluated with respect to the same iterate, whereas IAG uses different points for the first-order approximation and the proximity condition. We show that this critical difference leads to an incremental algorithm with a smaller linear convergence factor. Moreover, the linear convergence factor of the proposed DIAG method justifies the use of incremental methods to improve the performance of GD. In particular, we show

that the worse case scenario of DIAG is guaranteed to be better than the worse case scenario of GD. Based on our knowledge, this is the first incremental method which is guaranteed to improve the worse case performance of GD.

We start the paper by presenting the GD and IAG methods and studying their convergence guarantees for the case that the instantaneous functions f_i are strongly convex and their gradients ∇f_i are Lipschitz continuous (Section 2). We clarify the reason that the convergence analysis of IAG cannot guarantee the advantage of incremental methods with respect to GD. Then, we present the proposed DIAG method which uses both variable and gradient averages in oppose to IAG that only uses aggregated gradient average (Section 3). We explain the intuition behind this difference by comparing the function approximations used in these methods. Further, we suggest an efficient mechanism to implement the proposed DIAG algorithm that has the computational complexity of the order $O(p)$ which is significantly lower than of GD given by $\mathcal{O}(np)$ (Section 3.1). Further, we explain the connection between the proposed DIAG method and the majorization-minimization method (MISO) proposed in [17], and highlight the differences between these two algorithms (Remark 1).

The convergence analysis of the DIAG method is then presented (Section 4). We first prove a fundamental lemma that shows the error of DIAG at each iteration is strictly smaller than the average of the errors of the last n iterations (Lemma 1). We use this result to prove that the sequence of variables generated by DIAG converges to the optimal argument $\mathbf{x}^*$ (Proposition 2), and, in particular, the convergence rate of the iterates evaluated after each pass over the dataset is linear (Corollary 3). This linear convergence factor guarantees that one pass of DIAG is more efficient than one iteration of gradient descent, i.e., the upper bound for the error of DIAG after n gradient evaluations is strictly smaller than the one for GD. Then, we prove that the whole sequence of DIAG iterates is linearly convergent (Theorem 4) and characterize the linear convergence factor (Theorem 7). We extend our convergence results by studying the worst-case asymptotic rate of DIAG (Section 5). We use the Perron-Frobenius (PF) theory to show that an upper bound for the sequence of DIAG errors has an asymptotic linear convergence rate (Theorem 9). This analysis also matches the convergence results in Section 4.

We compare the performances of DIAG, GD, and IAG in solving a quadratic programming and a binary classification problem (Section 6). Numerical results for the quadratic programming confirm that DIAG outperforms GD. In particular, the relative performance of DIAG and GD does not vary by changing the problem condition number, while IAG is not preferable to GD when the problem condition number is relatively large. Moreover, DIAG outperforms IAG irrespective to the problem parameters. The convergence paths of these methods for the binary classification problem, which is a logistic regression minimization, conform the observations for the quadratic programming problem. Finally, we close the paper by concluding remarks (Section 7).

1.1. Notation. Vectors are written as $\mathbf{x} \in \mathbb{R}^p$ and matrices as $\mathbf{A} \in \mathbb{R}^{p\times p}$. Given n vectors $\mathbf{x}_i$, the vector $\mathbf{x} = [\mathbf{x}_1; \dots; \mathbf{x}_n]$ represents a stacking of the elements of each individual $\mathbf{x}_i$. We use $\|\mathbf{x}\|$ and $\|\mathbf{A}\|$ to denote the Euclidean norm of vector $\mathbf{x}$ and matrix $\mathbf{A}$, respectively. Given a function f its gradient $\mathbf{x}$ is denoted as $\nabla f(\mathbf{x})$.

2. Related Works and Preliminaries. Since the objective function in (1) is convex, descent methods can be used to find the optimal argument $\mathbf{x}^*$. In this paper, we are interested in studying methods that converge to the optimal argument of the global objective function f at a linear rate. It is customary for the linear convergence

analysis of first-order methods to assume that the functions are smooth and strongly convex. We formalize these conditions in the following assumption.

Assumption 1. *The functions f_i are differentiable and strongly convex with constant $\mu > 0$, i.e., for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$ we can write*

$$(\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y}))^T(\mathbf{x} - \mathbf{y}) \geq \mu \|\mathbf{x} - \mathbf{y}\|^2. \tag{2}$$

Moreover, the gradients ∇f_i are Lipschitz continuous with constant $L < \infty$, i.e., for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$ we have

$$\|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y})\| \leq L \|\mathbf{x} - \mathbf{y}\|. \tag{3}$$

The strong convexity of the functions f_i with constant μ implies that the global objective function f is also strongly convex with constant μ. Likewise, the Lipschitz continuity of the gradients ∇f_i with constant L yields Lipschitz continuity of the global objective function gradients ∇f with constant L. Note that the conditions in Assumption 1 are mild and hold for most large-scale machine learning applications such as, linear regression, logistic regression, least squares, and support vector machines.

The optimization problem in (1) can be solved using the gradient descent (GD) method [19]. The idea of GD is to update the current iterate $\mathbf{x}^k$ by descending through the negative direction of the current gradient $\nabla f(\mathbf{x}^k)$ with a proper stepsize ϵ^k. In other words, the update of GD for solving problem (1) at step k is defined as

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \epsilon^k \nabla f(\mathbf{x}^k) = \mathbf{x}^k - \frac{\epsilon^k}{n} \sum_{i=1}^{n} \nabla f_i(\mathbf{x}^k). \tag{4}$$

Convergence analysis of GD in [19] shows that the sequence of iterates $\mathbf{x}^k$ converges linearly to the optimal argument $\mathbf{x}^*$ if the stepsize is constant and satisfies $\epsilon^k = \epsilon < 2/L$. The fastest convergence rate is achieved by the stepsize $\epsilon = 2/(\mu + L)$ which leads to the linear convergence factor $(\kappa - 1)/(\kappa + 1)$, i.e.,

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \left(\frac{\kappa - 1}{\kappa + 1}\right)^k \|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{5}$$

where $\kappa := L/\mu$ is the global objective function condition number. Although, GD has a fast linear convergence rate, it is not computationally affordable in large-scale applications because of its high computational complexity. To comprehend this limitation, note that each iteration of GD requires n gradient evaluations which is not computationally affordable in large-scale applications with massive values of n. Stochastic gradient descent (SGD) arises as a natural solution in large-scale settings. SGD modifies the update of GD by approximating the gradient of the global objective function ∇f by the average of a small number of instantaneous gradients chosen uniformly at random from the set of n gradients. To be more precise, the update of SGD at step k is defined as

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \frac{\epsilon^k}{b} \sum_{i \in \mathcal{S}_b^k} \nabla f_i(\mathbf{x}^k), \tag{6}$$

where $\mathcal{S}_b^k$ is defined as a random set that contains the indices of b functions that are chosen for the update SGD at step k. Note that the components of the set $\mathcal{S}_b^k$ are

chosen uniformly at random from the set of indices $\{1, 2, \ldots, n\}$. Since the stochastic gradient $(1/b)\sum_{i\in\mathcal{S}_b^k}\nabla f_i(\mathbf{x}^k)$ is an unbiased estimator of the gradient $\nabla f(\mathbf{x}^k) = (1/n)\sum_{i=1}^n \nabla f_i(\mathbf{x}^k)$, the sequence of the iterates generated by SGD converges to the optimal argument in expectation. However, the convergence rate is sublinear and slower than the linear convergence of GD. In particular, the expected error $\|\mathbf{x}^k - \mathbf{x}^*\|^2$ of SGD is bounded above as $\mathbb{E}\left[\|\mathbf{x}^k - \mathbf{x}^*\|^2\right] \leq \mathcal{O}\left(1/k\right)$, for a diminishing stepsizes ϵ^k of the order $1/k$. It is worth mentioning that the expectation is taken with respect to the indices of the chosen random functions up to step k.

One may use a cyclic order instead of stochastic selection of functions in SGD which leads to the update of incremental gradient descent method (IGD) as in [3, 29]. Similar to the case for SGD, the sequence of iterates generated by the IGD method converges to the optimal argument at a sublinear rate of the order $\mathcal{O}\left(1/k\right)$ when the stepsize is diminishing. SGD and IGD are able to reduce the computational complexity of GD by requiring only one gradient evaluation per iteration; however, they both suffer from slow (sublinear) convergence rates.

The sublinear convergence rate of SGD has been improved recently by the stochastic average gradient method (SAG) which also can be interpreted as a stochastic incremental *aggregated* gradient method. The SAG method updates only one gradient per iteration and uses the average of the most recent version of all gradients – gradients of all functions $f_1, \ldots, f_n$ – as an approximation for the full gradient [23]. To be more specific, define $\mathbf{y}_i^k$ as the copy of the decision variable $\mathbf{x}$ for the last time that the function f_i's gradient is updated. Then, at each iteration, a random index i^k is chosen uniformly at random and the gradient of its corresponding function $\nabla f_{i^k}(\mathbf{x}^k)$ is evaluated and stored as $\nabla f_{i^k}(\mathbf{y}_i^k)$. Then, the update of SAG at step k is given by

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \frac{\epsilon}{n}\sum_{i=1}^n \nabla f_i(\mathbf{y}_i^k), \tag{7}$$

which uses the gradients of all the n functions evaluated at different time steps. The sequence of iterates generated by SAG converges linearly to $\mathbf{x}^*$ in expectation with respect to the choices of random indices, i.e.,

$$\mathbb{E}\left[\|\mathbf{x}^k - \mathbf{x}^*\|^2\right] \leq \left(1 - \frac{1}{8n\kappa}\right)^k C_0, \tag{8}$$

where C_0 is a constant independent of n and κ. However, the linear convergence constant of SAG in (8) is not necessarily better than the linear convergence constant of GD for a problem with comparable condition number. To be more precise, the residual $\|\mathbf{x}^k - \mathbf{x}^*\|$ of SAG decays by the factor of $(1 - 1/(8n\kappa))^n$ after a pass over the dataset which might not be better than the upper bound for the residual of GD that decays with the factor of $(\kappa - 1)/(\kappa + 1)$. As an example, for the problem that $n = \kappa = 100$, the worst case performance of GD after m passes over the set of functions (m iterations) is bounded above by $((\kappa - 1)/(\kappa + 1))^m\|\mathbf{x}^0 - \mathbf{x}^*\| \approx 0.98^m\|\mathbf{x}^0 - \mathbf{x}^*\|$, while the worst performance of SAG after m passes over the set of functions (mn iterations) is bounded above by $(1 - (1/8n\kappa))^{nm}C_0 \approx 0.9987^m C_0$. Note that the constants $\|\mathbf{x}^0 - \mathbf{x}^*\|$ and C_0 are negligible for sufficiently large m. Similar examples can be derived for other first-order stochastic methods with linear convergence rate in [9, 10, 13, 31, 26, 27, 32, 14, 17]. Beside this issue, the results for all these stochastic first-order methods hold in *expectation*. Thus, there is a positive probability that the sequence of iterates generated by these methods might not converge at a linear rate.

The other alternative for solving the optimization problem in (1) is the Incremental Aggregated Gradient (IAG) method which is a middle ground between GD and IGD. The IAG method requires one gradient evaluation per iteration, as in IG, while it approximates the gradient of the global objective function $\nabla f(\mathbf{x})$ by the average of the most recent gradient of all instantaneous functions [3], and it has a linear convergence rate, as in GD. In the IAG method, the functions are chosen in a cyclic order and it takes n iterations to have a pass over all the available functions. To introduce the update of IAG, recall the definition of $\mathbf{y}_i^k$ as the copy of the decision variable $\mathbf{x}$ for the last time that the function f_i's gradient is updated before step k. Then, the update of IAG is given by

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \frac{\epsilon}{n} \sum_{i=1}^{n} \nabla f_i(\mathbf{y}_i^k), \tag{9}$$

which is identical to the update of SAG in (7), and the only difference is in the scheme that the index i^k is chosen.

The convergence results in [29] provide global convergence and local linear convergence of IAG in a more general setting when each component function satisfies a local Lipschitzian error condition. More recently, a new convergence analysis of IAG has been studied in [11] which shows global linear convergence of IAG for strongly convex functions with Lipschitz continuous gradients. In particular, it has been shown that the sequence of iterates $\mathbf{x}^k$ generated by IAG satisfies the following inequality

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \left(1 - \frac{2}{25n(2n+1)(\kappa+1)^2}\right)^k \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{10}$$

Notice that the convergence rate of IAG is linear and eventually the error of IAG will be smaller than the errors of SGD and IGD which diminish with a sublinear rate of $\mathcal{O}(1/k)$. To compare the performance of GD and IAG it is fair to compare one iteration of GD with n iterations of IAG. This is reasonable since one iteration of GD requires n gradient evaluations, while IAG uses n gradient evaluations after n iterations. Comparing the decrement factors of GD in (5) and IAG after n gradient evaluations in (10) shows that there is no guarantee that IAG is preferable to GD for all choices of condition number κ and number of functions n, since we could face the scenario that

$$\left(\frac{\kappa-1}{\kappa+1}\right) < \left(1 - \frac{2}{25n(2n+1)(\kappa+1)^2}\right)^n. \tag{11}$$

As an example, for the problem with $n = \kappa = 100$, the inequality in (11) holds and the worst case performance of IAG is worse than the one for GD. Note that the bound for GD in (5) is strict and we can design a sequence which satisfies the equality case of the result in (5)[1]. However, the bound in (10) is not necessarily tight and it could be the reason that the comparison in (11) does not justify the use of IAG instead of GD. Our goal in this paper is to come up with a first-order incremental method that has a guaranteed upper bound which is better than the one for GD in (5). We propose this algorithm in the following section.

[1]Consider the quadratic programming $f(\mathbf{x}) = (1/2)\mathbf{x}^T\mathbf{A}\mathbf{x}$, where $\mathbf{A} = \text{diag}[\mu, L]$ which has the optimal argument $\mathbf{x}^* = \mathbf{0} \in \mathbb{R}^2$. Then, by setting $\epsilon = 2/(\mu + L)$ the sequence of iterates generated by GD satisfies the relation $\|\mathbf{x}^m - \mathbf{x}^*\| = \rho^m \|\mathbf{x}^0 - \mathbf{x}^*\|$.

3. Algorithm Definition. In this section, we propose a novel incremental gradient method that unlike other incremental methods is able to improve upon the worst case performance of GD. To do so, we first introduce a new interpretation of the IAG method. Recall the variable $\mathbf{y}_i^k$ as the copy of the decision variable $\mathbf{x}$ for the last time that function f_i is chosen for gradient update. The update of IAG in (9) can be interpreted as the solution of the optimization program

(12)
$$\mathbf{x}^{k+1} = \underset{\mathbf{x}\in\mathbb{R}^p}{\operatorname{argmin}} \left\{ \frac{1}{n}\sum_{i=1}^{n} f_i(\mathbf{y}_i^k) + \frac{1}{n}\sum_{i=1}^{n} \nabla f_i(\mathbf{y}_i^k)^T(\mathbf{x}-\mathbf{y}_i^k) + \frac{1}{n}\sum_{i=1}^{n} \frac{1}{2\epsilon}\|\mathbf{x}-\mathbf{x}^k\|^2 \right\}.$$

This interpretation shows that in the update of IAG each instantaneous function $f_i(\mathbf{x})$ is approximated by the following approximation

$$f_i(\mathbf{x}) \approx f_i(\mathbf{y}_i^k) + \nabla f_i(\mathbf{y}_i^k)^T(\mathbf{x}-\mathbf{y}_i^k) + \frac{1}{2\epsilon}\|\mathbf{x}-\mathbf{x}^k\|^2. \tag{13}$$

Notice that the first two terms $f_i(\mathbf{y}_i^k) + \nabla f_i(\mathbf{y}_i^k)^T(\mathbf{x}-\mathbf{y}_i^k)$ correspond to the first order approximation of the function f_i around the iterate $\mathbf{y}_i^k$. The last term which is $1/(2\epsilon)\|\mathbf{x}-\mathbf{x}^k\|^2$ is a proximal term that is added to the first order approximation. This approximation is different from the classic approximation that is used in first-order methods, since the first-order approximation is evaluated around the point $\mathbf{y}_i^k$ which is different from the iterate $\mathbf{x}^k$ used in the proximal term. This observation verifies that the IAG algorithm performs well when the delayed variables $\mathbf{y}_i^k$ are close to the current iterate $\mathbf{x}^k$ which is true when the stepsize ϵ is very small or the iterates are all close to the optimal solution.

We resolve this issue by introducing a different approach for approximating each component function f_i, In particular, we use the approximation

$$f_i(\mathbf{x}) \approx f_i(\mathbf{y}_i^k) + \nabla f_i(\mathbf{y}_i^k)^T(\mathbf{x}-\mathbf{y}_i^k) + \frac{1}{2\epsilon}\|\mathbf{x}-\mathbf{y}_i^k\|^2. \tag{14}$$

As we observe, the approximation in (14) is more consistent to classic first-order methods comparing to the one for IAG in (13). This is true since the first order approximation and the proximal term in (14) are evaluated with respect to the same point $\mathbf{y}_i^k$. Indeed, the approximation in (14) implies that the global objective function $f(\mathbf{x})$ can be approximated by

$$f(\mathbf{x}) \approx \frac{1}{n}\sum_{i=1}^{n} f_i(\mathbf{y}_i^k) + \frac{1}{n}\sum_{i=1}^{n} \nabla f_i(\mathbf{y}_i^k)^T(\mathbf{x}-\mathbf{y}_i^k) + \frac{1}{n}\sum_{i=1}^{n} \frac{1}{2\epsilon}\|\mathbf{x}-\mathbf{y}_i^k\|^2. \tag{15}$$

We can approximate the optimal argument of the global objective function f by minimizing its approximation in (15). Thus, the updated iterate $\mathbf{x}^{k+1}$ can be computed as the minimizer of the approximated global objective function in (15), i.e.,

(16)
$$\mathbf{x}^{k+1} = \underset{\mathbf{x}\in\mathbb{R}^p}{\operatorname{argmin}} \left\{ \frac{1}{n}\sum_{i=1}^{n} f_i(\mathbf{y}_i^k) + \frac{1}{n}\sum_{i=1}^{n} \nabla f_i(\mathbf{y}_i^k)^T(\mathbf{x}-\mathbf{y}_i^k) + \frac{1}{n}\sum_{i=1}^{n} \frac{1}{2\epsilon}\|\mathbf{x}-\mathbf{y}_i^k\|^2 \right\}.$$

Considering the convex programming in (16) we can derive a closed form expression for the update of $\mathbf{x}^{k+1}$ as

$$\mathbf{x}^{k+1} = \frac{1}{n}\sum_{i=1}^{n} \mathbf{y}_i^k - \frac{\epsilon}{n}\sum_{i=1}^{n} \nabla f_i(\mathbf{y}_i^k). \tag{17}$$

Algorithm 1 Double Incremental Aggregated Gradient method (DIAG)

Require: initial variables $\{\mathbf{y}_i^0\}_{i=1}^{i=n} = \mathbf{x}^0$ and gradients $\{\nabla f_i(\mathbf{y}_i^0)\}_{i=1}^{i=n}$
for $k = 0, 1, \dots$ **do**
  Compute the function index $i^k = \text{mod}(k, n) + 1$
  Compute $\mathbf{x}^{k+1} = \frac{1}{n}\sum_{i=1}^n \mathbf{y}_i^k - \frac{\epsilon}{n}\sum_{i=1}^n \nabla f_i(\mathbf{y}_i^k)$.
  Update the sum of variables $\sum_{i=1}^n \mathbf{y}_i^{k+1} = \mathbf{x}^{k+1} - \mathbf{y}_{i^k}^k + \sum_{i=1}^n \mathbf{y}_i^k$.
  Compute the gradient $\nabla f_{i^k}(\mathbf{x}^{k+1})$ and use it to update the sum of gradients $\sum_{i=1}^n \nabla f_i(\mathbf{y}_i^{k+1}) = \nabla f_{i^k}(\mathbf{x}^{k+1}) - \nabla f_{i^k}(\mathbf{y}_{i^k}^k) + \sum_{i=1}^n \nabla f_i(\mathbf{y}_i^k)$.
  Replace $\mathbf{y}_{i^k}^k$ and $\nabla f_{i^k}(\mathbf{y}_{i^k}^k)$ by $\mathbf{x}^{k+1}$ and $\nabla f_{i^k}(\mathbf{x}^{k+1})$, respectively. The other components remain unchanged, i.e., $\mathbf{y}_i^{k+1} = \mathbf{y}_i^k$ and $\nabla f_i(\mathbf{y}_i^{k+1}) = \nabla f_i(\mathbf{y}_i^k)$ for $i \neq i^k$.
end for

We refer to the proposed method with the update in (17) as the Double Incremental Aggregated Gradient method (DIAG). This appellation is justified considering that the update of DIAG requires the *incremented aggregate* of *both* variables and gradients and only uses *gradient* (first-order) information.

Notice that since we use a cyclic scheme, the set of variables $\{\mathbf{y}_1^k, \mathbf{y}_2^k, \dots, \mathbf{y}_n^k\}$ is equal to the set $\{\mathbf{x}^k, \mathbf{x}^{k-1}, \dots, \mathbf{x}^{k-n+1}\}$. Therefore, the iterate $\mathbf{x}^{k+1}$ is a function of the last n iterates $\{\mathbf{x}^k, \mathbf{x}^{k-1}, \dots, \mathbf{x}^{k-n+1}\}$. This observation has a fundamental role in the analysis of the proposed DIAG method – see Section 4.

Remark 1. *One may consider the proposed DIAG method as a cyclic version of the stochastic MISO algorithm in [17]. This is a valid interpretation; however, the convergence analyses and guarantees of these methods are quite different. The proposed DIAG method is designed based on the new interpretation in* (14) *that leads to a novel proof technique – see Lemma 1 – which is different from the analysis of MISO in [17]. This analytical difference leads to different convergence guarantees. In particular, the MISO algorithm cannot improve the performance of GD for all choices of n and κ, while the established theoretical results for DIAG in Section 4 guarantee that DIAG outperforms GD under any choices of n and κ.*

3.1. Implementation Details. Naive implementation of the update in (17) requires computation of sums of n vectors per iteration which is computationally costly. This unnecessary computation can be avoided by tracking the sums over time. To be more precise, the first sum in (17) which is the sum of the variables can be updated as

$$\sum_{i=1}^n \mathbf{y}_i^{k+1} = \mathbf{x}^{k+1} - \mathbf{y}_{i^k}^k + \sum_{i=1}^n \mathbf{y}_i^k, \tag{18}$$

where i^k is the index of the function that is chosen at step k. Likewise, the sum of gradients in (17) can be updated as

$$\sum_{i=1}^n \nabla f_i(\mathbf{y}_i^{k+1}) = \nabla f_{i^k}(\mathbf{x}^{k+1}) - \nabla f_{i^k}(\mathbf{y}_{i^k}^k) + \sum_{i=1}^n \nabla f_i(\mathbf{y}_i^k). \tag{19}$$

The proposed double incremental aggregated gradient (DIAG) method is summarized in Algorithm 1. The variables for all the copies of the vector $\mathbf{x}$ are initialized

by vector $\mathbf{x}^0$, i.e., $\mathbf{y}_1^0 = \cdots = \mathbf{y}_n^0 = \mathbf{x}^0$, and their corresponding gradients are stored in the memory. At each iteration k, the updated variable $\mathbf{x}^{k+1}$ is computed in Step 4 using the update in (17). The sums of variables and gradients are updated in Step 5 and 6, respectively, following the recursions in (18) and (19). In Step 7, the old variable $\mathbf{y}_{i^k}^k$ and gradient $\nabla f_{i^k}(\mathbf{y}_{i^k}^k)$ of the updated function f_{i^k} are replaced with their updated versions, i.e., $\mathbf{x}^{k+1}$ and $\nabla f_{i^k}(\mathbf{x}^{k+1})$, and the other components of the variables and gradients tables remain unchanged. In Step 3, the index i^k is updated in a cycling manner.

4. Convergence Analysis. In this section, we study the convergence properties of the proposed double incremental aggregated gradient method.

The following lemma characterizes an upper bound for the error $\|\mathbf{x}^{k+1} - \mathbf{x}^*\|$ in terms of the errors of the last n iterations.

LEMMA 1. *Consider the proposed double incremental aggregated gradient (DIAG) method in (17). If the conditions in Assumption 1 hold, and the stepsize ϵ is chosen as $\epsilon = 2/(\mu+L)$, the sequence of iterates $\mathbf{x}^k$ generated by DIAG satisfies the inequality*

$$\|\mathbf{x}^{k+1} - \mathbf{x}^*\| \leq \left(\frac{\kappa-1}{\kappa+1}\right)\left[\frac{\|\mathbf{x}^k - \mathbf{x}^*\| + \cdots + \|\mathbf{x}^{k-n+1} - \mathbf{x}^*\|}{n}\right], \tag{20}$$

where $\kappa = L/\mu$ is the objective function condition number.

Proof. See Appendix A. ∎

The result in Lemma 1 has a significant role in the analysis of DIAG. It shows that the error at step $k+1$ is smaller than the average of the last n errors where the decrement factor is the ratio $(\kappa-1)/(\kappa+1)$ which is strictly smaller than 1. The cyclic scheme is critical in proving the result in (20), since it allows to replace the sum $\sum_{i=1}^n \|\mathbf{y}_i^k - \mathbf{x}^*\|$ by the sum of the last n steps errors $\|\mathbf{x}^k - \mathbf{x}^*\| + \cdots + \|\mathbf{x}^{k-n} - \mathbf{x}^*\|$. Note that If we pick functions uniformly at random, as in MISO, it is not possible to write the expression in (20), even in expectation. We also cannot write an inequality similar to the one in (20) for the IAG method, although it uses a cyclic scheme. This contrast is originated by the difference that IAG only uses gradients average, whereas DIAG uses both variables and gradients averages. In the following theorem, we use the result in Lemma 1 to show that the sequence of variables $\mathbf{x}^k$ converges to the optimal argument $\mathbf{x}^*$.

PROPOSITION 2. *Consider the proposed double incremental aggregated gradient (DIAG) method in (17), and recall the definition $\rho := (\kappa-1)/(\kappa+1)$ where $\kappa = L/\mu$ is the problem condition number. If the conditions in Assumption 1 hold, and the stepsize ϵ is chosen as $\epsilon = 2/(\mu+L)$, then the residual $\|\mathbf{x}^k - \mathbf{x}^*\|$ of DIAG for iterations $k = 1, \ldots, n$ satisfies the inequality*

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \rho\left[1 - \frac{(k-1)(1-\rho)}{n}\right]\|\mathbf{x}^0 - \mathbf{x}^*\|, \quad \text{for } k = 1, \ldots, n, \tag{21}$$

and for the steps $k > n$ we have

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \rho^{\lfloor \frac{k-1}{n} \rfloor + 1}\left[1 - \frac{(1-\rho)}{n} \times \min\left\{1, \frac{n-1}{2}\right\}\right]\|\mathbf{x}^0 - \mathbf{x}^*\|, \quad \text{for } k > n, \tag{22}$$

where $\lfloor a \rfloor$ indicates the floor of a.

Proof. See Appendix B. ∎

The first outcome of the result in Proposition 2 is the convergence of the sequence $\|\mathbf{x}^k - \mathbf{x}^*\|$ to zero as k approaches infinity. The second result which we formalize in the following corollary shows that the sequence of error converges linearly after each pass over the dataset.

COROLLARY 3. *If the conditions in Proposition 2 are satisfied, the error of the proposed DIAG method after $m > 1$ passes over the set of functions f_i, which requires mn gradient evaluations and corresponds to the iterate $k = n(m-1)+1$, is bounded above by*

$$\|\mathbf{x}^{n(m-1)+1} - \mathbf{x}^*\| \leq \rho^m \left[1 - \frac{(1-\rho)}{n} \times \min\left\{1, \frac{n-1}{2}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{23}$$

Proof. Since we count the initial n gradient computations, the iterate $\mathbf{x}^1$ requires n gradient computations. After the first iteration each step requires only one gradient computation. Therefore, the total number of gradient computations to evaluate $\mathbf{x}^k$ is $n+k-1$. Conversely, the variable that has exactly used mn gradients to be evaluated is $\mathbf{x}^{n(m-1)+1}$. Thus, by setting $k = n(m-1)+1$ in (22) we obtain the residual of DIAG after mn passes over the set of functions and the claim in (23) follows. ∎

The result in Corollary 3 shows that the subsequence of the last iterates of each pass is linearly convergent. Moreover, the result in Corollary 3 verifies the advantage of DIAG method versus the full gradient descent (GD) method. In particular, it shows that the error of DIAG after $m > 1$ passes over the set of functions f_i corresponding to the iterate $k = n(m-1)+1$ is bounded above by $\rho^m[1 - ((1-\rho)/n) \times \min\{1, ((n-1)/2)\}]\|\mathbf{x}^0 - \mathbf{x}^*\|$ which is strictly smaller than the upper bound for the error of GD after m iterations (nm gradient computations) given by $\rho^m\|\mathbf{x}^0 - \mathbf{x}^*\|$. Therefore, the DIAG method outperforms GD for any choice of κ and $n > 1$; DIAG and GD are identical for $n = 1$.

Notice that after the first pass over the set of functions – iteration $k = 1$ for the DIAG method – the error of DIAG is upper bounded by $\rho\|\mathbf{x}^0 - \mathbf{x}^*\|$ based on the result in (21). This bound is identical to the result for GD after one pass over the set of functions, since the first iterations of GD and DIAG are identical.

Although the result in Corollary 3 implies that the DIAG method is preferable with respect to GD and shows linear convergence of a subsequence of iterates, it is not sufficient to prove linear convergence of the whole sequence of iterates generated by DIAG. To be more precise, the result in Corollary 3 shows that the subsequence of errors $\{\|\mathbf{x}^{kn} - \mathbf{x}^*\|\}_{k=0}^{\infty}$, which are associated with the variables at the end of each pass over the set of functions, is linearly convergent. However, we aim to show that the whole sequence $\{\|\mathbf{x}^k - \mathbf{x}^*\|\}_{k=0}^{\infty}$ is linearly convergent. To be more precise, our goal is to prove that the sequence of DIAG iterates satisfies $\|\mathbf{x}^k - \mathbf{x}^*\| \leq a\gamma^k\|\mathbf{x}^0 - \mathbf{x}^*\|$ for a constant $a > 0$ and a positive coefficient $0 \leq \gamma < 1$. In the following theorem, we show that this condition is satisfied for the DIAG method.

THEOREM 4. *Consider the introduced double incremental aggregated gradient (DIAG) method in (17). If the conditions in Assumption 1 hold, and the stepsize ϵ is chosen as $\epsilon = 2/(\mu + L)$, for $k \geq 1$ we can write*

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq a\gamma^k\|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{24}$$

if the constants $a > 0$ and $0 \leq \gamma < 1$ satisfy the following conditions

$$\rho\left(1 - \frac{(k-1)(1-\rho)}{n}\right) \leq a\gamma^k \quad \text{for} \quad k = 1, \ldots, n, \tag{25}$$

$$\gamma^{n+1} - \left(1 + \frac{\rho}{n}\right)\gamma^n + \frac{\rho}{n} \leq 0 \quad \text{for} \quad k > n. \tag{26}$$

Proof. See Appendix C. ∎

The result in Theorem 4 provides conditions on the constants a and γ such that the linear convergence inequality $\|\mathbf{x}^k - \mathbf{x}^*\| \leq a\gamma^k\|\mathbf{x}^0 - \mathbf{x}^*\|$ holds. However, it does not gurantee that the set of constants $\{a, \gamma\}$ that satisfy the required conditions in (25) and (26) is non-empty. In the following proposition we show that there exist constants a and γ satisfying these conditions.

PROPOSITION 5. *There exist constants $a > 0$ and $0 < \gamma < 1$ that satisfy the inequalities in (25) and (26). In other words, the set of feasible solutions for the system of inequalities in (25) and (26) is non-empty.*

Proof. See Appendix D. ∎

The result in Proposition 5 in conjunction with the result in Theorem 4 guarantees linear convergence of the iterates generated by the DIAG method. Although there are different pairs of $\{a, \gamma\}$ that satisfy the conditions in (25) and (26) and lead to the linear convergence result in (24), we are interested in finding the pair $\{a, \gamma\}$ that leads to the smallest linear convergence factor γ, i.e., the pair that guarantees the fastest linear convergence rate. To find the smallest γ, we should pick the smallest γ that satisfies the inequality $\gamma^{n+1} - (1 + \rho/n)\gamma^n + \rho/n \leq 0$. Then choose the smallest constant a that satisfies the conditions in (25) for the given γ. To do so, we first look at the properties of the function $h(\gamma) := \gamma^{n+1} - (1 + \rho/n)\gamma^n + \rho/n$ in the following lemma.

LEMMA 6. *Consider the function $h(\gamma) := \gamma^{n+1} - (1 + \rho/n)\gamma^n + \rho/n$ for $\gamma \in [0, 1)$. The function h has only one root γ_0 in the interval $[0, 1)$. Moreover, γ_0 is the smallest choice of γ that satisfies the condition in (26).*

Proof. The derivative of the function h is given by

$$\frac{d}{d\gamma}h = (n+1)\gamma^n - (n+\rho)\gamma^{n-1}. \tag{27}$$

Therefore, the only critical point of the function h in the interval $(0, 1)$ is $\gamma^* = (n + \rho)/(n + 1)$. The point γ^* is a local minimum for the function h, since the second derivative of the function h is positive at γ^*. Notice that the objective function value $h(\gamma^*) < 0$ is negative. Moreover, we know that $h(0) > 0$ and $h(1) = 0$. This observation shows that the function h has a root γ_0 between 0 and γ^* and this is the only root of function h in the interval $(0, 1)$. Thus, γ_0 is the smallest value of γ in the interval $(0, 1)$ that satisfies the condition in (26) ∎

The result in Lemma 6 shows that the unique root of the function $h(\gamma) := \gamma^{n+1} - (1 + \rho/n)\gamma^n + \rho/n$ in the interval $[0, 1)$ is the smallest γ that satisfies the condition in (26). We use this result to formalize the pair $\{a, \gamma\}$ with the smallest choice of γ which satisfies the conditions in (25) and (26).

THEOREM 7. *Consider the proposed double incremental aggregated gradient (DIAG) method in (17). Let the conditions in Assumption 1 hold, and set the stepsize as*

$\epsilon = 2/(\mu + L)$. Then, the sequence of iterates generated by DIAG is linearly convergent as

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq a_0 \gamma_0^k \|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{28}$$

where γ_0 is the unique root of the equation

$$\gamma^{n+1} - \left(1 + \frac{\rho}{n}\right)\gamma^n + \frac{\rho}{n} = 0, \tag{29}$$

in the interval $[0, 1)$ and a_0 is given by

$$a_0 = \max_{i \in \{1,\dots,n\}} \rho \left(1 - \frac{(i-1)(1-\rho)}{n}\right) \gamma_0^{-i}. \tag{30}$$

Proof. It follows from the results in Theorem 4 and Lemma 6. ∎

The result in Theorem 7 shows R-linear convergence of the DIAG iterates with the smallest linear convergence factor γ_0. In the following section, we show that a sequence which is an upper bound for the errors $\|\mathbf{x}^k - \mathbf{x}^*\|$ asymptotically converges Q-linearly to zero with the linear convergence constant in (29).

5. Worst-case asymptotic rate of DIAG. In this section, our aim is to provide an upper bound on the quantity $\|\mathbf{x}^k - \mathbf{x}^*\|$ which is the distance to optimality after k steps. It follows directly from (20) that the sequence d^k defined by the recursion

$$d^{k+1} = \rho \frac{d^k + d^{k-1} + \cdots + d^{k-n+1}}{n} \tag{31}$$

where $\rho = (\kappa - 1)/(\kappa + 1)$ and $d^j := \|\mathbf{x}^j - \mathbf{x}^*\|$ for $j = 0, 1, 2, \dots, n-1$ provides an upper bound for the sequence of the errors $\|\mathbf{x}^k - \mathbf{x}^*\|$ for all $k \geq 0$. In other words, we have the relation

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq d^k \quad \text{for all} \quad k \geq 0. \tag{32}$$

Next, we characterize the convergence properties of the sequence of d^k which provides an upper bound for the desired error sequence $\|\mathbf{x}^k - \mathbf{x}^*\|$. To do so, we rewrite the update of the sequence d^k in a matrix form which is more suitable for the analysis. Define the column vector $\mathbf{d}^k = [d^{k-1}; \dots; d^{k-n}] \in \mathbb{R}^n$ as the concatenation of the last n values of the sequence d^k up to step k. Considering the definition of $\mathbf{d}^k$ and the update of the sequence d^k we can write

$$\mathbf{d}^{k+1} = \mathbf{M}_\rho \mathbf{d}^k \quad \text{where} \quad \mathbf{M}_\rho := \begin{bmatrix} \frac{\rho}{n} & \frac{\rho}{n} & \dots & \frac{\rho}{n} \\ 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \tag{33}$$

We observe that the matrix $\mathbf{M}_\rho \in \mathbb{R}^{n \times n}$ is a non-negative matrix whose eigenvalues determine the asymptotic growth rate of the sequence $\mathbf{d}^k$ and hence of d^k. It is straightforward to check that the characteristic polynomial of $\mathbf{M}_\rho$ is

$$T(\lambda) = \lambda^n - \frac{\rho}{n}\lambda^{n-1} - \frac{\rho}{n}\lambda^{n-2} - \dots - \frac{\rho}{n} = \frac{\lambda^{n+1} - \left(1 + \frac{\rho}{n}\right)\lambda^n + \frac{\rho}{n}}{\lambda - 1} \tag{34}$$

whose roots are the eigenvalues of the matrix $\mathbf{M}_\rho$. In the remainder of this section, we will infer information about the eigenvalues of $\mathbf{M}_\rho$ using Perron-Frobenius (PF) theory. This theory is well developed for positive matrices where all the entries are strictly positive but $\mathbf{M}_\rho$ has zero entries and is therefore not positive. Nevertheless, the PF theory has been successfully extended to certain non-negative matrices called *irreducible* matrices. A square matrix $\mathbf{A}$ is called irreducible if for every i and j, there exists an r such that $\mathbf{A}^r(i,j) > 0$. In the next lemma, we prove that the matrix $\mathbf{M}_\rho$ is irreducible which will justify our use of PF theory developed for irreducible matrices.

LEMMA 8. *The matrix $\mathbf{M}_\rho$ is irreducible for any $\rho > 0$.*

Proof. By the definition of irreducibility, we need to show that for every i and j, there exists an r such that the r-th power of the matrix $\mathbf{M}_\rho$ is entrywise positive, i.e. $\mathbf{M}_\rho^r(i,j) > 0$. Let $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_n$ be the standard basis for $\mathbb{R}^n$. It suffices to show that we can choose $r = n$ for all i and j, i.e.

$$\mathbf{M}_\rho^n(i,j) = \mathbf{e}_i^T \mathbf{M}_\rho^n \mathbf{e}_j > 0, \quad \text{for all} \quad i, j. \tag{35}$$

By the definition of the recurrence (33), we have

$$\begin{bmatrix} d^{2n-1} \\ d^{2n-2} \\ \cdots \\ d^n \end{bmatrix} = \mathbf{d}^{2n} = (\mathbf{M}_\rho)^n \mathbf{d}^n = (\mathbf{M}_\rho)^n \begin{bmatrix} d^{n-1} \\ d^{n-2} \\ \cdots \\ d^0 \end{bmatrix}.$$

Fix any $j \in \{1, \ldots, n\}$ and choose $\mathbf{d}^n = \mathbf{e}_j$ (this would correspond to the initalization $d_{n-\ell} = 1$ for $\ell = j$ and $d_{n-\ell} = 0$ for $1 \le \ell \le n$ and $\ell \neq j$). Then, using the definition of $\mathbf{M}_\sigma$, it is easy to check that such an initialization of $\mathbf{d}^n$ leads to $d^n = \rho/n > 0$, $d^{n+1} > 0, \ldots, d^{2n-1} > 0$. Therefore, for every i and j, we have

$$(\mathbf{M}_\rho)^n(i,j) = \mathbf{e}_i^T (\mathbf{M}_\rho)^n \mathbf{e}_j = d^{2n-i} > 0 \tag{36}$$

which proves (35) and completes the proof. ∎

The following result shows that the sequence d^k converges to zero linearly with a constant γ_0 where γ_0 is defined by (29). We also derive upper and lower bounds on γ_0.

THEOREM 9. *Consider the constant $\rho = (\kappa - 1)/(\kappa + 1) \in (0, 1)$ and let $\lambda^*(\rho)$ be the spectral radius of the matrix $\mathbf{M}_\rho$. Then,*

(i) $\lambda^*(\rho)$ is the largest real root of the the polynomial characteristic polynomial $T(\lambda)$. Furthermore, it is a simple root.

(ii) We have the limit

$$\lim_{k \to \infty} d^{k+1}/d^k = \lambda^*(\rho). \tag{37}$$

(iii) For integer numbers $n > 1$ the constant $\lambda^*(\rho)$ is bounded below and above as

$$\rho \le \lambda^*(\rho) < \sqrt[n]{\rho}. \tag{38}$$

(iv) We have $\lambda^*(\rho) = \gamma_0$ where γ_0 is the largest real root of the polynomial $h(\lambda) := \lambda^{n+1} - (1 + (\rho/n))\lambda^n + (\rho/n)$ in the interval $[0, 1)$.

Proof. See Appendix E. ∎

The results in Theorem 9 study the convergence properties of the sequence d^k defined in (31) which is an upper bound for the sequence of DIAG error $\|\mathbf{x}^k - \mathbf{x}^*\|$. The last result in Theorem 9 shows that the constant $\lambda^*(\rho)$ which is the spectral radius of the matrix $\mathbf{M}_\rho$, is equal to the linear convergence factor γ_0 of DAIG defined as the root of the polynomial in (29). The second result indicates that the sequence d^k has an asymptotic linear convergence rate with the constant $\lambda^*(\rho) = \gamma_0$. This convergence rate matches the convergence analysis of the proposed DIAG method in Theorem 7 which shows that the smallest guaranteed linear convergence factor for the DIAG method is γ_0. The second result and the upper bound for the constant $\lambda^*(\rho)$ in the third result together verify that DIAG improves GD. To be more precise, they show that the linear convergence factor of an upper bound for the error of DIAG method is smaller than $\sqrt[n]{\rho}$ asymptotically. Thus, n iterations of DIAG is better than one iteration of GD with decrement factor ρ. The lower bound for $\lambda^*(\rho)$ in (38) shows that linear convergence factor of DIAG cannot be better than the one for GD and the equality holds when $n = 1$.

6. Numerical experiments. In this section, we study the performance of the proposed DIAG method and compare it with existing alternative first-order methods. To do so, we first apply DIAG to solve a family of quadratic programming problems. Then, we evaluate the performance of DIAG and other first-order methods in solving a logistic regression minimization problem.

6.1. Quadratic programming example. To study the effect of number of functions n and problem condition number κ on the performance of the GD, IAG, and DIAG methods, we first apply these algorithms in solving a quadratic programming problem, where we can tune the problem condition number. In particular, consider the optimization problem

$$\min_{\mathbf{x}\in\mathbb{R}^p} f(\mathbf{x}) := \frac{1}{n}\sum_{i=1}^{n} \frac{1}{2}\mathbf{x}^T\mathbf{A}_i\mathbf{x} + \mathbf{b}_i^T\mathbf{x}, \tag{39}$$

where each matrix $\mathbf{A}_i \in \mathbb{R}^{p\times p}$ is a positive definite diagonal matrix and each vector $\mathbf{b}_i \in \mathbb{R}^p$ is randomly chosen from the box $[0, 1]^p$. To control the problem condition number, the first $p/2$ diagonal elements of $\mathbf{A}_i$ are chosen uniformly at random from the interval $[1, 10^1, \ldots, 10^{\eta/2}]$ and its last $p/2$ elements chosen from the interval $[1, 10^{-1}, \ldots, 10^{-\eta/2}]$. This selection resulting in the sum matrix $\sum_{i=1}^n \mathbf{A}_i$ having eigenvalues in the range $[n10^{-\eta/2}, n10^{\eta/2}]$. In our experiments, we fix the variable dimension as $p = 20$ and the number of functions as $n = 200$. Moreover, the stepsizes of GD and DIAG are set as their best theoretical stepsizes which are $\epsilon_{GD} = 2/(\mu+L)$ and $\epsilon_{DIAG} = 2/(\mu + L)$, respectively. Note that the stepsize suggested in [11] for IAG is $\epsilon_{IAG} = 0.32/((nL)(L + \mu))$; however, this choice of stepsize is very slow in practice. Thus, we use the stepsize $\epsilon_{IAG} = 2/(nL)$ which performs better than the one suggested in [11].

To have a fair comparison between these three methods, we compare them in terms of the total number of gradient evaluations. Note that comparing these methods in terms of the total number of iterations would not be fair since each iteration of GD requires n gradient evaluations, while IAG and DIAG only require one gradient computation per iteration.

We first consider the case that $\eta = 1$ and use the realization with condition number $\kappa = 10$ to have a relatively small condition number. Fig. 1 demonstrates

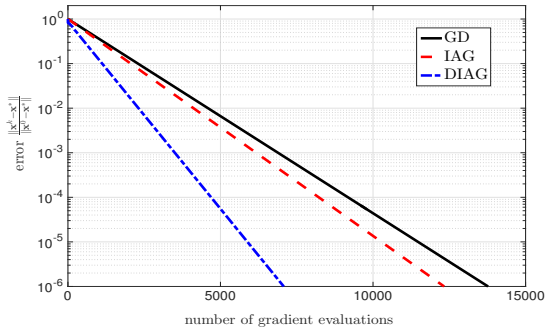

FIG. 1. *Relative error of GD, IAG, and DIAG versus number of gradient evaluations for the quadratic programming in (39) with $n = 200$ and $\kappa = 10$. The IAG method performs slightly better than GD, while DIAG has the fastest convergence path.*

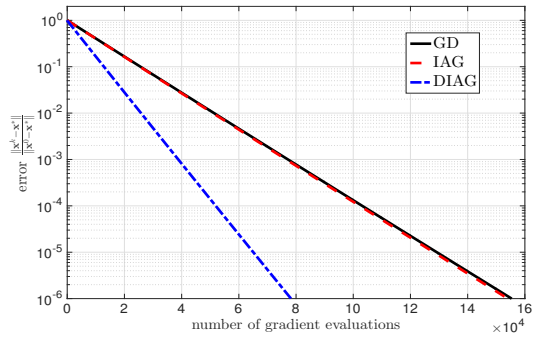

FIG. 2. *Relative error of GD, IAG, and DIAG versus number of gradient evaluations for the quadratic programming in (39) with $n = 200$ and $\kappa = 117$. In this setting, IAG and GD have similar convergence paths, and the best performance belongs to DIAG.*

the convergence paths of the normalized error $\|\mathbf{x}^k - \mathbf{x}^*\|/\|\mathbf{x}^0 - \mathbf{x}^*\|$ for IAG, DIAG, and GD when $n = 200$ and $\kappa = 10$. As we observe, IAG performs slightly better than GD, while the best performance belongs to DIAG. To be more precise, DIAG requires $7,069$ gradient evaluations (approximately 35 passes over the dataset) to achieve the relative error of $\|\mathbf{x}^k - \mathbf{x}^*\|/\|\mathbf{x}^0 - \mathbf{x}^*\| = 10^{-6}$, while IAG requires $12,330$ gradient evaluations (approximately 61 passes over the dataset) to achieve the same accuracy. The GD method has the worst performance and achieves the relative error $\|\mathbf{x}^k - \mathbf{x}^*\|/\|\mathbf{x}^0 - \mathbf{x}^*\| = 10^{-6}$ after 68 iterations which is equivalent to $13,600$ gradient evaluations.

Comparison of the convergence guarantees for GD, IAG, and DIAG shows that the IAG method is more sensitive to the problem condition number, since the linear convergence factor of IAG is of the order $1 - \mathcal{O}(1/\kappa^2)$, while the linear convergence factor of GD and DIAG are at the order of $1 - \mathcal{O}(1/\kappa)$. To study the effect of problem condition number in practice, we increase the constant η to have a poorly conditioned problem. In particular, we increase the problem condition number by setting $\eta = 2$ and using the realization with condition number $\kappa = 117$. Fig. 2 illustrates the performance of these methods for the case that $n = 200$ and $\kappa = 117$. We observe that the convergence path of IAG is almost identical to the one for GD. This observation verifies that the performance of IAG worsens more significantly by increasing the problem condition number. Interestingly, the relative performance of

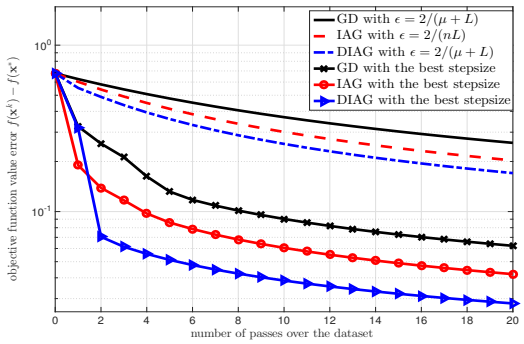
GD with $\epsilon = 2/(\mu + L)$
IAG with $\epsilon = 2/(nL)$
DIAG with $\epsilon = 2/(\mu + L)$
GD with the best stepsize
IAG with the best stepsize
DIAG with the best stepsize
objective function value error $f(\mathbf{x}^k) - f(\mathbf{x}^*)$
number of passes over the dataset

FIG. 3. *Convergence paths of GD, IAG, and DIAG for the binary classification application.*

DIAG and GD does not change by increasing the problem condition number. To be more specific, for the case that $n = 200$ and $\kappa = 117$, GD and IAG reach the relative error $\|\mathbf{x}^k - \mathbf{x}^*\|/\|\mathbf{x}^0 - \mathbf{x}^*\| = 10^{-6}$ after 1.54×10^5 gradient evaluations, while DIAG requires only 7.8×10^4 gradient computations to achieve the same accuracy.

6.2. Logistic Regression minimization. In this section, we compare the performance of GD, IAG, and DIAG in solving a binary classification problem. Consider the given training set $\mathcal{S} = \{\mathbf{u}_i, l_i\}_{i=1}^{i=n}$ which contains n realizations of the feature vectors $\mathbf{u}_i \in \mathbb{R}^p$ and respective label l_i where the labels are either -1 or 1. The goal is to find the optimal classifier $\mathbf{x}^* \in \mathbb{R}^p$ that minimizes the regularized logistic loss which is given by

$$\min_{\mathbf{x}\in\mathbb{R}^p} f(\mathbf{x}) := \frac{1}{n}\sum_{i=1}^{n} \log(1 + \exp(-l_i \mathbf{x}^T \mathbf{u}_i)) + \frac{\lambda}{2}\|\mathbf{x}\|^2, \tag{40}$$

where the regularization term $(\lambda/2)\|\mathbf{x}\|^2$ is added to avoid overfitting. The problem in (40) is a particular case of the problem in (1) when the function f_i is defined as $f_i(\mathbf{x}) = \log(1 + \exp(-l_i \mathbf{x}^T \mathbf{u}_i)) + (\lambda/2)\|\mathbf{x}\|^2$.

Note that the objective function f in (40) is strongly convex with the constant $\mu = \lambda$ and its gradients are Lipschitz continuous with the constant $L = \lambda + \zeta/4$ where $\zeta = \max_i \mathbf{u}_i^T \mathbf{u}_i$. It is easy to verify that the instantaneous functions f_i are also strongly convex with constant $\mu = L$, and their gradients are Lipschitz continuous with constant $L = \lambda + \zeta/4$. This observation shows that condition in Assumption 1 hold for the logistic regression problem in (40).

We apply GD, IAG, and DIAG to solve the logistic regression problem in (40) for the MNIST dataset [15]. We only use the samples that correspond to digits 0 and 8 and assign label $l_i = 1$ to the samples that correspond to digit 8 and label $l_i = -1$ to those associated with digit 0. We get a total of $n = 11,774$ training examples, each of dimension $p = 784$. The objective function error $f(\mathbf{x}^k) - f(\mathbf{x}^*)$ of the GD, IAG, and DIAG methods versus the number of passes over the dataset are shown in Fig. 3. We report the results for two different sets of stepsizes. We first use the stepsizes $\epsilon_{GD} = 2/(\mu + L)$, $\epsilon_{IAG} = 2/(nL)$, and $\epsilon_{DIAG} = 2/(\mu + L)$ as in the quadratic programming. Then, we hand-optimize the choice of stepsize for each of these three methods separately and report the best performance. We observe that, for both theoretical and hand-optimized stepsizes, the proposed DIAG method outperforms GD and IAG.

7. Conclusion. In this paper we proposed a novel incremental method for solving the average of a set of n smooth and strongly convex functions. The proposed double incremental aggregated gradient method (DIAG) uses the aggregated average of both variables and gradients to update its iterate in oppose to classic cyclic incremental methods that only use the average of gradients. The convergence analysis of the DIAG method guarantees improvement with respect to the gradient descent (GD) method. This result makes DIAG the first cyclic incremental method that improves GD under all circumstances. Moreover, we showed that the sequence of iterates generated by DIAG is linearly convergent. Numerical experiments matched the theoretical results and showcased the advantage of DIAG relative to GD and the classic incremental aggregated gradient method (IAG).

Appendix A. Proof of Lemma 1.

Consider the update in (17). Subtract the optimal argument $\mathbf{x}^*$ from both sides of the equality to obtain

$$\mathbf{x}^{k+1} - \mathbf{x}^* = \frac{1}{n}\sum_{i=1}^{n}(\mathbf{y}_i^k - \mathbf{x}^*) - \frac{\epsilon}{n}\sum_{i=1}^{n}\nabla f_i(\mathbf{y}_i^k). \tag{41}$$

Note that the global objective function gradient at the optimal point is null, i.e., $(1/n)\sum_{i=1}^{n}\nabla f_i(\mathbf{x}^*) = \mathbf{0}$. This observation in conjunction with the expression in (41) leads to

$$\begin{aligned}\mathbf{x}^{k+1} - \mathbf{x}^* &= \frac{1}{n}\sum_{i=1}^{n}\left(\mathbf{y}_i^k - \mathbf{x}^*\right) - \frac{\epsilon}{n}\sum_{i=1}^{n}\left(\nabla f_i(\mathbf{y}_i^k) - \nabla f_i(\mathbf{x}^*)\right)\\ &= \frac{1}{n}\sum_{i=1}^{n}\Big[\mathbf{y}_i^k - \mathbf{x}^* - \epsilon\left(\nabla f_i(\mathbf{y}_i^k) - \nabla f_i(\mathbf{x}^*)\right)\Big]. \end{aligned}\tag{42}$$

Compute the norm of both sides in (42), and use the Cauchy-Schwarz inequality to obtain

$$\|\mathbf{x}^{k+1} - \mathbf{x}^*\| \leq \frac{1}{n}\sum_{i=1}^{n}\left\|\mathbf{y}_i^k - \mathbf{x}^* - \epsilon\left(\nabla f_i(\mathbf{y}_i^k) - \nabla f_i(\mathbf{x}^*)\right)\right\|. \tag{43}$$

Now we proceed to derive an upper bound for each summand in (43). According to the result in [[24]; page 13], if the functions f_i are μ-strongly convex and their gradients are L-Lipschitz continuous we can show that

$$\left\|\mathbf{y}_i^k - \mathbf{x}^* - \epsilon\left(\nabla f_i(\mathbf{y}_i^k) - \nabla f_i(\mathbf{x}^*)\right)\right\| \leq \max\{|1-\epsilon\mu|, |1-\epsilon L|\}\|\mathbf{y}_i^k - \mathbf{x}^*\|. \tag{44}$$

By setting the stepsize ϵ in (44) as $\epsilon = 2/(\mu + L)$, we can write

$$\left\|\mathbf{y}_i^k - \mathbf{x}^* - \epsilon\left(\nabla f_i(\mathbf{y}_i^k) - \nabla f_i(\mathbf{x}^*)\right)\right\| \leq \frac{\kappa - 1}{\kappa + 1}\,\|\mathbf{y}_i^k - \mathbf{x}^*\|, \tag{45}$$

where $\kappa = L/\mu$ is the function f_i condition number. By replacing the summands in the right hand side of (43) with their upper bounds $((\kappa-1)/(\kappa+1))\|\mathbf{y}_i^k - \mathbf{x}^*\|$, as shown in (45), we can show that the residual $\|\mathbf{x}^{k+1} - \mathbf{x}^*\|$ is bounded above as

$$\|\mathbf{x}^{k+1} - \mathbf{x}^*\| \leq \left(\frac{\kappa - 1}{\kappa + 1}\right)\sum_{i=1}^{n}\frac{\|\mathbf{y}_i^k - \mathbf{x}^*\|}{n} \tag{46}$$

Note that in the DIAG method we use a cyclic scheme to update the variables. Thus, the set of variables $\{\mathbf{y}_1^k, \ldots, \mathbf{y}_n^k\}$ is identical to the set of the last n iterates before the iterate $\mathbf{x}^{k+1}$ which is given by $\{\mathbf{x}^k, \ldots, \mathbf{x}^{k-n+1}\}$. Thus, we can replace the sum in $\sum_{i=1}^n \|\mathbf{y}_i^k - \mathbf{x}^*\|$ in (46) by the sum $\sum_{i=1}^n \|\mathbf{x}^{k-i+1} - \mathbf{x}^*\|$ and the claim in (20) follows.

Appendix B. Proof of Proposition 2.

Consider the definition of the constant $\rho := (\kappa - 1)/(\kappa + 1)$ where $\kappa = L/\mu$ is the objective function condition number. Thus, if all the copies $\mathbf{y}_i$ are initialized at $\mathbf{x}^0$, the result in Lemma 1 implies that

$$\|\mathbf{x}^1 - \mathbf{x}^*\| \leq \rho \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{47}$$

We can use the same inequality for the second iterate to obtain

$$\|\mathbf{x}^2 - \mathbf{x}^*\| \leq \frac{\rho}{n} \|\mathbf{x}^1 - \mathbf{x}^*\| + \frac{\rho(n-1)}{n} \|\mathbf{x}^0 - \mathbf{x}^*\| \tag{48}$$

Replace $\|\mathbf{x}^1 - \mathbf{x}^*\|$ in (48) by its upper bound in (47) and regroup the terms to obtain

$$\begin{aligned} \|\mathbf{x}^2 - \mathbf{x}^*\| &\leq \frac{\rho^2}{n} \|\mathbf{x}^0 - \mathbf{x}^*\| + \frac{\rho(n-1)}{n} \|\mathbf{x}^0 - \mathbf{x}^*\| \\ &= \rho \left[1 - \frac{1-\rho}{n}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|. \end{aligned} \tag{49}$$

Repeat the same process for the third residual $\|\mathbf{x}^3 - \mathbf{x}^*\|$ to obtain

$$\begin{aligned} \|\mathbf{x}^3 - \mathbf{x}^*\| &\leq \frac{\rho}{n} \|\mathbf{x}^2 - \mathbf{x}^*\| + \frac{\rho}{n} \|\mathbf{x}^1 - \mathbf{x}^*\| + \frac{\rho(n-2)}{n} \|\mathbf{x}^0 - \mathbf{x}^*\| \\ &\leq \rho \left[1 - \frac{2(1-\rho)}{n} - \frac{\rho(1-\rho)}{n^2}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \end{aligned} \tag{50}$$

where in the second inequality we use the bounds in (47) and (48). Since the term $-\rho(1-\rho)/n^2$ is negative we can drop this term and show that the residual $\|\mathbf{x}^3 - \mathbf{x}^*\|$ is upper bounded by

$$\|\mathbf{x}^3 - \mathbf{x}^*\| \leq \rho \left[1 - \frac{2(1-\rho)}{n}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{51}$$

By following the same logic we can show that for the first n residuals $\{\|\mathbf{x}^k - \mathbf{x}^*\|\}_{k=1}^n$ the following inequality holds

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \rho \left[1 - \frac{(k-1)(1-\rho)}{n}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \quad \text{for } k = 1, \ldots, n. \tag{52}$$

Thus, the result in (21) holds.

Now we proceed to show that the result in (22) hold for $k = n+1, \ldots, 2n$. According to the result in Lemma 1 we can write

$$\|\mathbf{x}^{n+1} - \mathbf{x}^*\| \leq \frac{\rho}{n} \|\mathbf{x}^n - \mathbf{x}^*\| + \cdots + \frac{\rho}{n} \|\mathbf{x}^1 - \mathbf{x}^*\| \tag{53}$$

By replacing each summand in the right hand side of (53) by its upper bound in (52) we obtain

$$\begin{aligned} \|\mathbf{x}^{n+1} - \mathbf{x}^*\| &\leq \frac{\rho}{n} \left[\sum_{i=1}^n \rho \left[1 - \frac{(i-1)(1-\rho)}{n}\right]\right] \|\mathbf{x}^0 - \mathbf{x}^*\| \\ &= \rho^2 \left[1 - \frac{(1-\rho)(n-1)}{2n}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|. \end{aligned} \tag{54}$$

It follows from the result in (54) that the residual $\|\mathbf{x}^{n+1} - \mathbf{x}^*\|$ is upper bounded by

$$\|\mathbf{x}^{n+1} - \mathbf{x}^*\| \leq \rho^2 \left[1 - \frac{1-\rho}{n} \times \min\left\{1, \frac{(n-1)}{2n}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{55}$$

Therefore, the result in (22) holds for $k = n+1$. To prove the claim for $k = n+2$, first note that based on the upper bounds in (52) we can show that

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \rho \left[1 - \frac{(1-\rho)}{n}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \quad \text{for } k = 2, \ldots, n. \tag{56}$$

Considering the following inequality

$$\left[1 - \frac{1-\rho}{n}\right] \leq \left[1 - \frac{1-\rho}{n} \times \min\left\{1, \frac{(n-1)}{2n}\right\}\right], \tag{57}$$

and the upper bounds in (56) we obtain that

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \rho \left[1 - \frac{1-\rho}{n} \times \min\left\{1, \frac{(n-1)}{2n}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \quad \text{for } k = 2, \ldots, n. \tag{58}$$

In addition, the result in (55) and the fact that $\rho < 1$ imply that the term $\|\mathbf{x}^{n+1} - \mathbf{x}^*\|$ is also can be upper bounded by

$$\|\mathbf{x}^{n+1} - \mathbf{x}^*\| \leq \rho \left[1 - \frac{1-\rho}{n} \times \min\left\{1, \frac{(n-1)}{2n}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{59}$$

Now, based on the result in Lemma 1, we can write

$$\|\mathbf{x}^{n+2} - \mathbf{x}^*\| \leq \frac{\rho}{n} \left[\|\mathbf{x}^{n+1} - \mathbf{x}^*\| + \cdots + \|\mathbf{x}^2 - \mathbf{x}^*\|\right]. \tag{60}$$

The inequalities in (58) and (59) show that all the summands in (60) are bounded by the same upper bound. Replace these term by the upper bound to obtain

$$\|\mathbf{x}^{n+2} - \mathbf{x}^*\| \leq \rho^2 \left[1 - \frac{1-\rho}{n} \times \min\left\{1, \frac{(n-1)}{2n}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{61}$$

and the claim in (22) for $k = n+2$ follows.

Since the constant ρ is strictly less than 1, we can replace the upper bound in (61) by

$$\|\mathbf{x}^{n+2} - \mathbf{x}^*\| \leq \rho \left[1 - \frac{1-\rho}{n} \times \min\left\{1, \frac{(n-1)}{2n}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{62}$$

Using the same argument, we can say that the upper bound in (62) holds for $k = n+2, \ldots, 3$ and apply the result in Lemma 1 to show that the distance $\|\mathbf{x}^{n+3} - \mathbf{x}^*\|$ is bounded above by

$$\|\mathbf{x}^{n+3} - \mathbf{x}^*\| \leq \rho^2 \left[1 - \frac{1-\rho}{n} \times \min\left\{1, \frac{(n-1)}{2n}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{63}$$

which yields the claim in (22) for $k = n+3$. By repeating the steps in (62) and (63) we can conclude that the result in (22) holds for $k = n+1, \ldots, 2n$.

The proof for steps $k > 2n$ is similar to the argument used for the steps $k = n+1, \ldots, 2n$, although we write it in a formal manner by using induction.

Assume that for for $k = nj+1, \ldots, nj+n$ the following inequality holds

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \rho^{\lfloor \frac{k-1}{n} \rfloor + 1} \left[1 - \frac{(1-\rho)}{n} \times \min\left\{1, \frac{n-1}{2}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{64}$$

We intend to prove the same inequalities hold for $k = n(j+1)+1, \ldots, n(j+1)+n$. Note that the result in (64) is satisfied for $j = 1$, which corresponds to the iterates $k = n+1, \ldots, 2n$, and the base of induction holds.

As we assume that the result in (64) holds for $k = nj+1, \ldots, nj+n$, we obtain that

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \rho^{j+1} \left[1 - \frac{(1-\rho)}{n} \times \min\left\{1, \frac{n-1}{2}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{65}$$

for $k = nj+1, \ldots, nj+n$. According to Lemma 1, the residual $\|\mathbf{x}^{n(j+1)+1} - \mathbf{x}^*\|$ is bounded above by

$$\|\mathbf{x}^{n(j+1)+1} - \mathbf{x}^*\| \leq \frac{\rho}{n} \|\mathbf{x}^{nj+1} - \mathbf{x}^*\| + \cdots + \frac{\rho}{n} \|\mathbf{x}^{nj+n} - \mathbf{x}^*\|. \tag{66}$$

Replacing the summands in the right hand side of (66) by their upper bound in (65) implies that

$$\|\mathbf{x}^{n(j+1)+1} - \mathbf{x}^*\| \leq \rho^{j+2} \left[1 - \frac{(1-\rho)}{n} \times \min\left\{1, \frac{n-1}{2}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{67}$$

which yields the claim in (64) for $k = n(j+1)+1$.

Now use the result in (67) and the inequality $\rho < 1$ to obtain

$$\|\mathbf{x}^{n(j+1)+1} - \mathbf{x}^*\| \leq \rho^{j+1} \left[1 - \frac{(1-\rho)}{n} \times \min\left\{1, \frac{n-1}{2}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{68}$$

Proceed by writing the result in Lemma 1 for $k = n(j+1)+2$ to obtain

$$\|\mathbf{x}^{n(j+1)+2} - \mathbf{x}^*\| \leq \frac{\rho}{n} \|\mathbf{x}^{nj+2} - \mathbf{x}^*\| + \cdots + \frac{\rho}{n} \|\mathbf{x}^{nj+n+1} - \mathbf{x}^*\|. \tag{69}$$

Replace the summands in the right hand side of (69) by their upper bounds in (65) and (68), which are the same upper bounds, to obtain

$$\|\mathbf{x}^{n(j+1)+2} - \mathbf{x}^*\| \leq \rho^{j+2} \left[1 - \frac{(1-\rho)}{n} \times \min\left\{1, \frac{n-1}{2}\right\}\right] \|\mathbf{x}^0 - \mathbf{x}^*\|, \tag{70}$$

and the claim in (64) for $k = n(j+1)+2$ follows. By repeating the steps from (68) to (70) we can show that the same result holds for $k = k = n(j+1)+3, \ldots, k = n(j+1)+n$. Thus, the inequality in (64) holds for $k = n(j+1)+1, \ldots, n(j+1)+n$. The induction is complete which implies that the claim in (22) holds.

Appendix C. Proof of Theorem 4.

Considering the result in (21) and the definition of the constant a in (25) we obtain that

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq a\gamma^k \|\mathbf{x}^0 - \mathbf{x}^*\|, \qquad \text{for} \quad k = 1, \ldots, n. \tag{71}$$

Thus, the inequality in (24) holds for steps $k = 1, \ldots, n$.

Now we proceed to show that the claim in (24) also holds for $k > n$. To do so, we use an induction argument. Let's assume we aim to show that the inequality in (24) holds for $k = j$, while it holds for the last n iterates $k = j-1, \ldots, j-n$. According to the result in Lemma 1 we can write

$$\|\mathbf{x}^j - \mathbf{x}^*\| \leq \rho \left[\frac{\|\mathbf{x}^{j-1} - \mathbf{x}^*\| + \cdots + \|\mathbf{x}^{j-n} - \mathbf{x}^*\|}{n} \right], \tag{72}$$

where $\rho = (\kappa - 1)/(\kappa + 1)$. Based on the induction assumption, for steps $k = j - 1, \ldots, j-n$, the result in (24) holds. Thus, we can replace the terms in the right hand side of (72) by the upper bounds from (24). This substitution implies

$$\begin{aligned} \|\mathbf{x}^j - \mathbf{x}^*\| &\leq \frac{\rho a}{n} \left[\gamma^{j-1} + \cdots + \gamma^{j-n} \right] \|\mathbf{x}^0 - \mathbf{x}^*\| \\ &= \frac{\rho a \gamma^{j-n}(1-\gamma^n)}{n(1-\gamma)} \|\mathbf{x}^0 - \mathbf{x}^*\| \end{aligned} \tag{73}$$

Rearranging the terms in (26) allows us to show that $(\rho(1-\gamma^n))/(n(1-\gamma))$ is bounded above by γ^n. This is true since

$$\begin{aligned} \gamma^{n+1} - \left(1 + \frac{\rho}{n}\right)\gamma^n + \frac{\rho}{n} \leq 0 &\iff \frac{\rho}{n}(1-\gamma^n) - \gamma^n(1-\gamma) \leq 0 \\ &\iff \frac{\rho(1-\gamma^n)}{n(1-\gamma)} \leq \gamma^n. \end{aligned} \tag{74}$$

Therefore, we can replace the term $(\rho(1-\gamma^n))/(n(1-\gamma))$ in (73) by its upper bound γ^n to obtain

$$\|\mathbf{x}^j - \mathbf{x}^*\| \leq a\gamma^j \|\mathbf{x}^0 - \mathbf{x}^*\|. \tag{75}$$

The result in (75) completes the proof. Thus, by induction the claim in (24) holds for all $k \geq 1$ if the conditions in (25) and (26) satisfied.

Appendix D. Proof of Proposition 5.

To prove the claim in Proposition 5 we first derive the following lemma.

LEMMA 10. *For all $n \geq 1$ and $0 \leq \phi \leq 1$ we have*

$$\left(1 - \frac{\phi}{n}\right)^n \leq \left(1 - \frac{\phi}{n+1}\right)^{n+1} \tag{76}$$

Proof. Consider the function $h(x) = (1-(\phi/x))^x$ for $x > 1$. The natural logarithm of the function $h(x)$ is given by $\ln(h(x)) = x \ln(1 - (\phi/x))$. Compute the derivative of both sides with respect to x to obtain

$$\frac{dh}{dx} \times \frac{1}{h(x)} = \ln\left(1 - \frac{\phi}{x}\right) + x \times \frac{\frac{\phi}{x^2}}{1 - \frac{\phi}{x}} \tag{77}$$

By multiplying both sides by $h(x)$, replacing $h(x)$ by the expression $\left(1 - \frac{\phi}{x}\right)^x$, and simplifying the terms we obtain that the derivative of the function $h(x)$ is given by

$$\frac{dh}{dx} = \left(1 - \frac{\phi}{x}\right)^x \left[\ln\left(1 - \frac{\phi}{x}\right) + \frac{\frac{\phi}{x}}{1 - \frac{\phi}{x}} \right]. \tag{78}$$

Note that the sum $\ln(1-u)+u/(1-u)$ is always positive for $0<u<1$. By setting $u:=\phi/x$, we can conclude that the term in the right hand side of (78) is positive for $x>1$. Therefore, the derivative dh/dx is always positive for $x>1$. Thus, the function $h(x)$ is an increasing function for $x>1$ and we can write

$$\left(1-\frac{\phi}{n}\right)^n \leq \left(1-\frac{\phi}{n+1}\right)^{n+1}, \tag{79}$$

for $0\leq\phi\leq 1$ and $n>1$. It remains to show that the same claim is also valid for $n=1$ which is equivalent to the inequality

$$1-\phi \leq \left(1-\frac{\phi}{2}\right)^2. \tag{80}$$

It is trivial to show (80) holds, and, therefore, the claim in (76) holds for all $n\geq 1$. □

Now proceed to prove the claim in Proposition 5 using the result in Lemma 10. To prove that the feasible set of the condition in (26) is non-empty we show that $\gamma=\rho^{1/n}$ satisfies the inequality in (26). In other words,

$$\rho^{\frac{n+1}{n}} - \left(1+\frac{\rho}{n}\right)\rho + \frac{\rho}{n} \leq 0 \tag{81}$$

Divide both sides of (81) by ρ and regroupe the terms to obtain the following ineqaulity

$$\rho \leq \left(1-\frac{1-\rho}{n}\right)^n, \tag{82}$$

which is equivalent to (81). In other words, the inequality in (82) is a necessary and sufficient condition for the condition in (81).

Recall the result in Lemma 10. By setting $\phi=1-\rho$ we obtain that

$$\rho = \left(1-\frac{1-\rho}{1}\right)^1 \leq \left(1-\frac{1-\rho}{2}\right)^2 \leq \cdots \leq \left(1-\frac{1-\rho}{n}\right)^n, \tag{83}$$

for $n\geq 1$. Thus, the inequality in (82) holds, and, consequently, the inequality in (81) is valid. Therefore, $\gamma=\rho^{1/n}$ satisfies the inequality in (26).

Then, we can define a as the smallest constant that satisfies (25) for the choice $\gamma=\rho^{1/n}$, which is given by

$$a = \max_{k=1,\ldots,n}\left(1-\frac{(k-1)(1-\rho)}{n}\right)\rho^{1-\frac{k}{n}}. \tag{84}$$

Therefore, $\gamma=\rho^{1/n}$ and the constant a in (84) satisfy the conditions in (25) and (26), and the claim in Proposition 5 follows.

Appendix E. Proof of Theorem 9.

(*i*) This follows directly from the Perron-Frobenius theorem for irreducible non-negative matrices [12, Theorem 8.4.4].

(*ii*) By [12, Theorem 8.5.1], we also have

$$\lim_{k\to\infty}\frac{\mathbf{M}_\rho^k}{\lambda^*(\rho)^k} = \mathbf{u}\mathbf{v}^T \tag{85}$$

where $\mathbf{u}$, $\mathbf{v}$ are the right and left eigenvectors of $\mathbf{M}_\rho$ corresponding to the eigenvalue $\lambda^*(\rho)$ normalized to satisfy $\mathbf{v}^T\mathbf{u} = 1$. Note also that

$$d^k = \mathbf{e}_1^T\mathbf{d}^{k+1} = \mathbf{e}_1^T\mathbf{M}_\rho^{k-n+1}\mathbf{d}^n \tag{86}$$

Therefore,

$$\lim_{k\to\infty} d^{k+1}/d^k = \lim_{k\to\infty} \frac{\mathbf{e}_1^T\mathbf{M}_\rho^{k-n+2}\mathbf{d}^n}{\mathbf{e}_1^T\mathbf{M}_\rho^{k-n+1}\mathbf{d}^n} = \lim_{k\to\infty} \frac{\mathbf{e}_1^T\mathbf{M}_\rho^{k-n+1}\mathbf{d}^n/\lambda^*(\rho)^{k-n+1}}{\mathbf{e}_1^T\mathbf{M}_\rho^{k-n}\mathbf{d}^n/\lambda^*(\rho)^{k-n+1}} \tag{87}$$

where in the last inequality we divided both numerator and denominator with the same factor $\lambda^*(\rho)^{k-n+1}$. Simplifying and using (85), we obtain

$$\lim_{k\to\infty} d^{k+1}/d^k = \lim_{k\to\infty} \lambda^*(\rho)\frac{\mathbf{e}_1^T\mathbf{u}\mathbf{v}^T\mathbf{d}^n}{\mathbf{e}_1^T\mathbf{u}\mathbf{v}^T\mathbf{d}^n} = \lambda^*(\rho). \tag{88}$$

(*iii*) A direct consequence of [Theorem 8.1.22][12] is that $\lambda^*(\rho) \geq \rho$, this proves the lower bound on $\lambda^*(\rho)$. To get the upper bound, let $\mathbf{1} = [1\ 1\ \ 1 \dots 1]^T$ be the vector of ones. We will show that

$$\mathbf{M}_\rho^{2n}\mathbf{1} < \rho^2\mathbf{1} \tag{89}$$

where the notation " $<$ " denotes the componentwise inequality for vectors. Then, by [12, Corollary 8.1.29], this would imply

$$\lambda^*(\rho)^{2n} < \rho^2$$

which is equivalent to the desired upper bound. It is a straightforward computation to show that if we set $\mathbf{d}^n = \mathbf{1}$, then after a simple induction argument we obtain $d^n = \rho$ and $d^{n+1} < \rho$, $d^{n+2} < \rho$, $\dots$, $d^{2n-1} < \rho$, i.e.

$$\mathbf{d}^{2n} = \begin{bmatrix} d^{2n-1} \\ d^{2n-2} \\ \dots \\ d^n \end{bmatrix} = \mathbf{M}_\rho^n\mathbf{d}^n = \mathbf{M}_\rho^n\mathbf{1} = \rho\mathbf{v} \tag{90}$$

for a vector $\mathbf{v} = [v_1, v_2, \dots, v_n]^T$, where $v_i < 1$ if $i < n$ and $v_n = 1$. Using similar arguments we can write

$$\mathbf{d}^{3n} = \begin{bmatrix} d^{3n-1} \\ d^{3n-2} \\ \dots \\ d^{2n} \end{bmatrix} = \mathbf{M}_\rho^{2n}\mathbf{d}^n = \mathbf{M}_\rho\mathbf{M}_\rho^n\mathbf{1} = \rho\mathbf{M}_\rho\mathbf{v} \tag{91}$$

and a straightforward computation shows that $\mathbf{M}_\rho\mathbf{v} < \rho\mathbf{1}$. Combining this inequality with the previous equation proves (89) and concludes the proof.

(*iv*) It follows from (34) that the roots of the polynomial h is the same as the roots of the polynomial T except that h has an additional root at 1. By part (*i*) and (*iii*), $\lambda^*(\rho)$ is the largest real root of T and $0 < \lambda^*(\rho) < 1$. Therefore, it is also the largest real root of the function h over the interval $(0, 1)$ which is equal to γ_0.

REFERENCES

[1] D. P. Bertsekas, *A new class of incremental gradient methods for least squares problems*, SIAM Journal on Optimization, 7 (1997), pp. 913–926.

[2] D. P. Bertsekas, *Incremental proximal methods for large scale convex optimization*, Mathematical Programming, 129 (2011), pp. 163–195.

[3] D. Blatt, A. O. Hero, and H. Gauchman, *A convergent incremental gradient method with a constant step size*, SIAM Journal on Optimization, 18 (2007), pp. 29–51.

[4] L. Bottou, *Large-scale machine learning with stochastic gradient descent*, In Proceedings of COMPSTAT'2010, (2010), pp. 177–186.

[5] L. Bottou and Y. L. Cun, *On-line learning for very large datasets*, in Applied Stochastic Models in Business and Industry, vol. 21, pp. 137-151, 2005.

[6] F. Bullo, J. Cortes, and S. Martinez, *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*, Princeton University Press, 2009.

[7] Y. Cao, W. Yu, W. Ren, and G. Chen, *An overview of recent progress in the study of distributed multi-agent coordination*, IEEE Transactions on Industrial Informatics, 9 (2013), pp. 427–438.

[8] V. Cevher, S. Becker, and M. Schmidt, *Convex optimization for big data: Scalable, randomized, and parallel algorithms for big data analytics*, IEEE Signal Processing Magazine, 31 (2014), pp. 32–43.

[9] A. Defazio, F. R. Bach, and S. Lacoste-Julien, *SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives*, in Advances in Neural Information Processing Systems 27, Montreal, Quebec, Canada, 2014, pp. 1646–1654.

[10] A. Defazio, J. Domke, and T. S. Caetano, *Finito: A faster, permutable incremental gradient method for big data problems*, in Proceedings of the 31th International Conference on Machine Learning, ICML, Beijing, China, 2014, pp. 1125–1133.

[11] M. Gürbüzbalaban, A. Ozdaglar, and P. Parrilo, *On the convergence rate of incremental aggregated gradient algorithms*, arXiv preprint arXiv:1506.02081, (2015).

[12] R. A. Horn and C. R. Johnson, *Topics in matrix analysis*, Cambridge UP, New York, 1991.

[13] R. Johnson and T. Zhang, *Accelerating stochastic gradient descent using predictive variance reduction*, in Advances in Neural Information Processing Systems 26, Lake Tahoe, Nevada, United States, 2013, pp. 315–323.

[14] J. Konečnỳ and P. Richtárik, *Semi-stochastic gradient descent methods*, arXiv preprint arXiv:1312.1666, (2013).

[15] Y. LeCun, C. Cortes, and C. J. Burges, *The MNIST database of handwritten digits*, 1998.

[16] C. G. Lopes and A. H. Sayed, *Diffusion least-mean squares over adaptive networks: Formulation and performance analysis*, IEEE Transactions on Signal Processing, 56 (2008), pp. 3122–3136.

[17] J. Mairal, *Incremental majorization-minimization optimization with application to large-scale machine learning*, SIAM Journal on Optimization, 25 (2015), pp. 829–855.

[18] A. Mokhtari and A. Ribeiro, *DSA: decentralized double stochastic averaging gradient algorithm*, Journal of Machine Learning Research, 17 (2016), pp. 1–35.

[19] Y. Nesterov, *Introductory lectures on convex optimization*, vol. 87, Springer Science & Business Media, 2004.

[20] A. Ribeiro, *Ergodic stochastic optimization algorithms for wireless communication and networking*, IEEE Transactions on Signal Processing, 58 (2010), pp. 6369–6386.

[21] A. Ribeiro, *Optimal resource allocation in wireless communication and networking*, EURASIP Journal on Wireless Communications and Networking, 2012 (2012), pp. 1–19.

[22] H. Robbins and S. Monro, *A stochastic approximation method*, The Annals of Mathematical Statistics, (1951), pp. 400–407.

[23] N. L. Roux, M. W. Schmidt, and F. R. Bach, *A stochastic gradient method with an exponential convergence rate for finite training sets*, in Advances in Neural Information Processing Systems 25, Lake Tahoe, Nevada, United States., 2012, pp. 2672–2680.

[24] E. K. Ryu and S. Boyd, *Primer on monotone operator methods*, Appl. Comput. Math, 15 (2016), pp. 3–43.

[25] S. Shalev-Shwartz and N. Srebro, *SVM optimization: inverse dependence on training set size*, in Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, 2008, pp. 928–935.

[26] S. Shalev-Shwartz and T. Zhang, *Stochastic dual coordinate ascent methods for regularized loss*, The Journal of Machine Learning Research, 14 (2013), pp. 567–599.

[27] S. Shalev-Shwartz and T. Zhang, *Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization*, Mathematical Programming, 155 (2016), pp. 105–145.

[28] P. Tseng, *An incremental gradient (-projection) method with momentum term and adaptive stepsize rule*, SIAM Journal on Optimization, 8 (1998), pp. 506–531.
[29] P. Tseng and S. Yun, *Incrementally updated gradient methods for constrained and regularized optimization*, Journal of Optimization Theory and Applications, 160 (2014), pp. 832–853.
[30] N. D. Vanli, M. Gurbuzbalaban, and A. Ozdaglar, *Global convergence rate of proximal incremental aggregated gradient methods*, arXiv preprint arXiv:1608.01713, (2016).
[31] L. Xiao and T. Zhang, *A proximal stochastic gradient method with progressive variance reduction*, SIAM Journal on Optimization, 24 (2014), pp. 2057–2075.
[32] L. Zhang, M. Mahdavi, and R. Jin, *Linear convergence with condition number independent access of full gradients*, in Advances in Neural Information Processing Systems 26, Lake Tahoe, Nevada, United States., 2013, pp. 980–988.